6

Repeated Games, Trigger Strategies, and Tacit Collusion

An economic transaction is a solved political problem. Economics has gained the title of Queen of the Social Sciences by choosing solved political problems as its domain.

Lerner (1972):259

Money confers short-side power.

Bowles and Gintis (1993b)

6.1 Introduction

When a game is repeated an indefinite number of times by the same players, the anomalies associated with finitely repeated games (§5.11) disappear. Moreover such games are of great practical importance—we go to work every day for the same boss, we wake up every morning with the same spouse, we buy the same products, the same firms compete in the same markets, and so on. Except, in each case, sometimes we don't—which is why the repeated game is so strategically important.

The central result in the theory of repeated games is the Folk Theorem, which is an essentially negative result in classical game theory saying that repeated games have zillions equilibria. Of course, the Folk Theorem has virtually no standing in evolutionary game theory, because few of the equilibria it supports are stable with respect to a plausible evolutionary dynamic (I say this from intuition—the topic has not been systematically studied). Nevertheless it is worth presenting this theorem in some detail, since it exemplifies the logic of repeated games. I also present *Tacit Collusion* (§6.3) in class, because of its practical importance in industrial regulation and economic policy. *Tit-for-Tat* (§6.11) is also a classroom topic, because it is analytically challenging and historically important.

The chapter closes with an extended analysis of *Contingent Renewal*

VISIT…

Contracts (§6.14), which summarizes some of my own research with Samuel Bowles, expanding on Bowles and Gintis (1993b).

6.2 Big Fish and Little Fish

Many species of fish are attacked by parasites that attach themselves to gills and inner mouth parts. Often such fish will form a symbiotic relationship with a smaller species of fish for whom the parasite is a major food source. Mutual trust is involved, however, since the larger fish must avoid the temptation of eating the smaller fish, and the smaller fish must avoid the temptation of taking a chunk out of the larger fish, thereby obtaining a meal with much less work than picking around for the tiny parasites. This scenario, which is doubtless even more common and more important for humans than for fish, is explored in the following problem.

	C	D
C	5,5	-3,8
D	8,-3	0,0

Suppose Big Fish and Little Fish play the prisoner's dilemma listed on the right at times $t = 0, 1, 2, \ldots$. The payoff to each is the sum of the payoffs over all periods, weighted by discount factor δ, with $0 < \delta < 1$—a *discount factor* δ relates to a *discount rate* ρ by the formula $\delta = 1/(1+\rho)$. We call the game played in each period the *stage game* of a *repeated game* in which at each period the players can condition their moves on the complete previous history of the various stages. A strategy that dictates following one course of action until a certain condition is met and then following a different strategy for the rest of the game is called a *trigger strategy*.

Theorem 6.1. *The cooperative solution (5,5) can be achieved as a subgame perfect Nash equilibrium of the repeated game if δ is sufficiently close to unity and each player uses the trigger strategy of cooperating as long as the other player cooperates, and defecting forever if the other player defects on one round.*

Proof: We use the fact that for any discount factor δ with $0 < \delta < 1$,

$$1 + \delta + \delta^2 + \ldots = \frac{1}{1-\delta}.$$

To see this, write

$$\begin{aligned} x &= 1 + \delta + \delta^2 + \ldots \\ &= 1 + \delta(1 + \delta + \delta^2 + \ldots) \\ &= 1 + \delta x, \end{aligned}$$

from which the result follows.[1]

Now suppose both agents play the trigger strategy. Then the payoff to each is $5/(1-\delta)$. Suppose a player uses another strategy. This must involve cooperating for a number (possibly zero) of periods, then defecting forever; for once the player defects, his opponent will defect forever, the best response to which is to defect forever. Consider the game from the time t at which the first player defects. We can call this $t = 0$ without loss of generality. A fish that defects receives 8 immediately and nothing thereafter. Thus the cooperate strategy is Nash if and only if $5/(1-\delta) \geq 8$, or $\delta \geq 3/8$. When δ satisfies this inequality, the pair of trigger strategies are also subgame perfect, since the situation in which both parties defect forever is Nash subgame perfect. Q.E.D.

This gives us an elegant solution to the problem, but in fact there are lots of other subgame perfect Nash equilibria to this game. For instance, Big Fish and Little Fish can trade off trying to bite and eat each other:

Theorem 6.2. *Consider the following trigger strategy for Little Fish: alternate $C, D, C, \ldots$ as long as Big Fish alternates $D, C, D, \ldots$. If Big Fish deviates from this pattern, defect forever. Suppose Big Fish plays the complementary strategy: alternate $D, C, D, \ldots$ as long as Little Fish alternates $C, D, C, \ldots$. If Little Fish deviates from this pattern, defect forever. These two strategies form a subgame perfect Nash equilibrium for δ sufficiently close to unity.*

Proof: The payoffs are $-3, 8, -3, 8, \ldots$ for Little Fish and $8, -3, 8, -3, \ldots$ for Big Fish. Let x be the payoffs to Little Fish. Little Fish gets -3 today, 8 in the next period and then gets to play the game all over again starting two periods from today. Thus $x = -3 + 8\delta + \delta^2 x$. Solving this, we get

$$x = \frac{8\delta - 3}{1 - \delta^2}.$$

The alternative for Little Fish to defect at some point, the most advantageous being when it is his turn to get -3. He then gets zero in that and all future periods. Thus cooperating is Nash if and only if $x \geq 0$, which is equivalent to $8\delta - 3 \geq 0$, or $\delta \geq 3/8$. Q.E.D.

[1] There is a faster way of arriving at the same result. Consider a repeated game that pays 1 now and in each period to a certain player, and the discount factor is δ. Let x be the value of the game to the player. The player receives 1 now and then gets to play exactly the same game in the next period. Since the value of the game in the next period is x, its present value is δx. Thus $x = 1 + \delta x$, so $x = 1/(1-\delta)$.

The above equilibrium is pretty silly, because both fish do worse than by fully cooperating, which they are capable of doing exactly when they can sustain this equilibrium. The value of the game to Big Fish is x, and it is easy to check that this is always less than $5/(1-\delta)$. Little Fish does even worse. Despite being silly, at least this pair of strategies resonates with us—it is similar to the spiteful quarrels that family members often manage to sustain over long periods of time. But there are many other subgame perfect Nash equilibria that simply defy any semblance of reality. Consider, for instance:

Theorem 6.3. *Suppose both Little Fish and Big Fish alternate $C, D, C, \ldots$ as long as the other does the same, so the payoffs are $5, 0, 5, 0, \ldots$, both players otherwise using the trigger strategy of defecting forever. This pair of strategies forms a subgame perfect Nash equilibrium for suitably low discount rates.*

Proof: The present value of the stream $5, 0, 5, 0, \ldots$ is the solution x to $x = 5 + 0\delta + x\delta^2$, which gives

$$x = \frac{5}{1-\delta^2}.$$

By defecting instead of cooperating, a player gets 8 immediately and zero forever after, for a total payoff of 8. Thus cooperating is Nash if and only if $x \geq 8$, which gives $\delta \geq \sqrt{3/8} \approx 0.61$.

Can you think of other creative subgame perfect Nash equilibria? Here is one: Little Fish plays C except on her birthday and any time the time period t that is a multiple of 7777; Big Fish plays C except on Greek Orthodox (fish) holidays. Each tolerates the other's play, but responds to a deviation by defecting forever. Here is another: Little Fish defects and Big Fish cooperates on the first round, then cooperates while Big Fish defects for just enough rounds that Little Fish's gain is zero (the last round might need a mixed strategy on the part of Big Fish). Then we repeat it all over again, with Big Fish defecting on the first round, then cooperating and Little Fish defecting until any gain of Big Fish's is wiped out. And so on, *ad infinitum*. Thus the game is worth exactly zero to each. But it is Nash subgame perfect!

6.3 Tacit Collusion

Consider a duopoly operating over an infinite number of periods $t = 1, 2, \ldots$. Suppose the duopolists are price-setters, so a pure strategy for

firms 1 and 2 in period t are prices $p_1^t, p_2^t \geq 0$, respectively, conditioned on the history of prices in previous time periods, and a pure strategy for the whole game is a sequence of strategies, one for each period t. Suppose the profits in period t are given by $\pi_1(p_1^t, p_2^t)$ for firm 1 and $\pi_2(p_1^t, p_2^t)$ for firm 2. The payoff to the firms for the whole game are then

$$\pi_1 = \sum_{t=0}^{\infty} \delta^t \pi_1(p_1^t, p_2^t), \qquad \pi_2 = \sum_{t=0}^{\infty} \delta^t \pi_2(p_1^t, p_2^t).$$

where δ is the common discount factor for the firms.

To specify the function $\pi_i(p_1^t, p_2^t)$, suppose the two firms have no fixed cost and constant marginal cost c, the firms face a downward-sloping demand curve, and the lowest price producer gets the whole market. Also, if the two producers have the same price, they share the market equally.

Theorem 6.4. *There is a subgame perfect Nash equilibrium of this game in which $p_1^t = p_2^t = c$ for $t = 1, 2, \ldots$.*

Notice that this is the 'competitive' equilibrium in which profits are zero and price equals marginal cost. The existence of this Nash equilibrium is called *Bertrand's Paradox* because it seems implausible (though hardly paradoxical!) that two firms in a duopoly actually behave this way.

Theorem 6.5. *Suppose $\delta > 50\%$, p^m is the monopoly price (i.e., the price that maximizes the profits of a monopolist), and $c \leq p \leq p^m$. The pair of strategies in which $p_1^t = p$ for firm 1 and $p_2^t = p$ for firm 2, $t = 1, 2, \ldots$, and each firm responds to a deviation from this strategy on the part of the other player by setting price equal to c forever, form a subgame perfect Nash equilibrium.*

We call a Nash equilibrium of this type *tacit collusion*.

Proof: Choose p satisfying the conditions of the theorem. Let $\pi(p)$ be total industry profits if price p is charged by both firms, so $\pi(p) \geq 0$. Suppose firm 2 follows the specified strategy. The payoff to firm 1 for following this strategy is

$$\pi^1(p) = \frac{\pi(p)}{2}\left(1 + \delta + \delta^2 + \ldots\right) = \frac{\pi(p)}{2(1-\delta)}. \tag{6.1}$$

The payoff to firm 1 for defecting on the first round by charging an amount $\epsilon > 0$ less than p is $\pi^1(p - \epsilon)$. Thus $2(1 - \delta) < 1$ is sufficient

for Nash. Clearly the strategy is subgame perfect, since the Bertrand solution in which each firm charges marginal cost is Nash. Q.E.D.

Intuition tells us that tacit collusion is more difficult to sustain when there are many firms. The following theorem, the proof of which we leave to the reader (just replace the '2' in the denominators of equation (6.1) by 'n') shows that this is correct.

Theorem 6.6. *Suppose there are $n > 1$ firms in the industry, but the conditions of supply and demand remain as before, the set of firms with the lowest price in a given time period sharing the market equally. Then if $\delta > 1 - 1/n$ and $c \leq p \leq p^m$, the trigger strategies in which each firm sets price equal to p in each period, and each firm responds to a deviation from this strategy on the part of another firm by setting price equal to c forever, form a subgame perfect Nash equilibrium.*

Another market condition that reduces the likelihood that tacit collusion can be sustained is incomplete knowledge on the part of the colluding firms. For instance, we have

Theorem 6.7. *Suppose there is an n-firm oligopoly, as described above, but a firm that has been defected upon cannot implement the trigger strategy until $k > 1$ periods have passed. Then tacit collusion can hold only if $\delta^{k+1} > 1 - 1/n$.*

This theorem shows that *tacit collusion is less likely to hold, the more easily a firm can 'hide' its defection.* This leads to the counterintuitive, but quite correct, conclusion that making contractual conditions public knowledge ("putting all your cards on the table") may *reduce* rather than *increase* the degree of competition in an industry.

Proof: If it takes k periods after a defection to retaliate, the gain from defection is

$$\pi(p)(1 + \delta + \delta^2 + \ldots + \delta^k) = \frac{\pi(p)}{1 - \delta}\left(1 - \delta^{k+1}\right),$$

from which the result immediately follows, since the present value of cooperating forever is $\pi(p)/n(1 - \delta)$. Q.E.D.

Theorem 6.8. *Suppose in each period $t = 1, 2, \ldots$ industry demand is either low, $q = D(p)$, or high, $q = D(p) + \alpha$ where $\alpha > 0$, with equal probability. Then under conditions of monopoly, the price p_h^m in the high demand period exceeds the price p_l^m in the low demand period. However*

under some conditions a duopoly engaging in tacit collusion to maximize their total profits may charge **lower** *prices in high demand period than in the low demand period.*

Proof: Suppose a monopolist maximizes $\pi = (p-c)(D(p)+\beta)$, giving first order condition $\pi_p = D(p) + \beta + (p-c)D'(p) = 0$. Differentiating this totally with respect to β gives

$$\pi_{pp}\frac{dp}{d\beta} + \pi_{p\beta} = 0.$$

However $\pi_{pp} < 0$ by the second order condition, and $\pi_{p\beta} = 1$, so $dp/d\beta > 0$. This shows that $p_h^m > p_l^m$ (let β go from zero to α).

The payoff to the trigger strategy in which each firm chooses the monopoly price in each period is

$$V = \sum_{t=0}^{\infty} \delta^t \left[\frac{1}{2}\frac{\pi_l^m}{2} + \frac{1}{2}\frac{\pi_h^m}{2}\right] = \frac{\pi_l^m + \pi_h^m}{4(1-\delta)}.$$

When is this a Nash equilibrium? Clearly if a firm does not benefit from defecting in the high demand period, it also does not benefit by defecting in the low demand period. So we need only find out when defecting in the high demand period is not profitable. The additional payoff from defecting in the high demand period is $\pi_h^m/2$, and the cost is δV (since the firm loses V, but starting in the next period), so for tacit collusion to be Nash, we must have $\delta \geq \pi_h^m/2V$, which simplifies to

$$\delta \geq \frac{2\pi_h}{3\pi_h^m + \pi_l^m}. \tag{6.2}$$

It is easy to see that the right hand side of this expression lies in the interval (1/2,2/3). Thus tacit collusion is harder in this case than in the case where demand does not fluctuate.

Suppose however that δ lies in the interval (1/2,2/3) but (6.2) is *violated*. Then it may be possible to maintain tacit collusion by charging *less* than the monopoly price in one or both periods, thus lowering the temptation to defect. Let $V(p_l, p_h)$ be the return to each firm when both firms charge p_l in the low demand state and p_h in the high demand state. Let us choose p_l and p_h to maximize $V(p_l, p_h)$ subject to the constraints that ensure that tacit collusion is Nash. If π_l and π_h are profits in the low and high demand states, we must have

$$\frac{\pi_s}{2} \leq \delta V(p_l, p_h) \qquad \text{for } s = l, h. \tag{6.3}$$

We shall show that this requires setting $p_l = p_l^m$ and setting

$$\pi_h = a\pi_l^m,$$

where $a = \frac{\delta}{2-3\delta}$, provided $\pi_h \leq \pi_h^m$. To see this, we rewrite the constraints (6.3) as

$$\pi_s \leq \delta \frac{\pi_l + \pi_h}{2(1-\delta)} \qquad \text{for } s = l, h,$$

which reduce to

$$\pi_l \leq a\pi_h, \qquad \pi_h \leq a\pi_l.$$

Note that $a \geq 1$, since $1/2 < \delta < 2/3$. The two possible cases are exhibited in Figure 6.1.

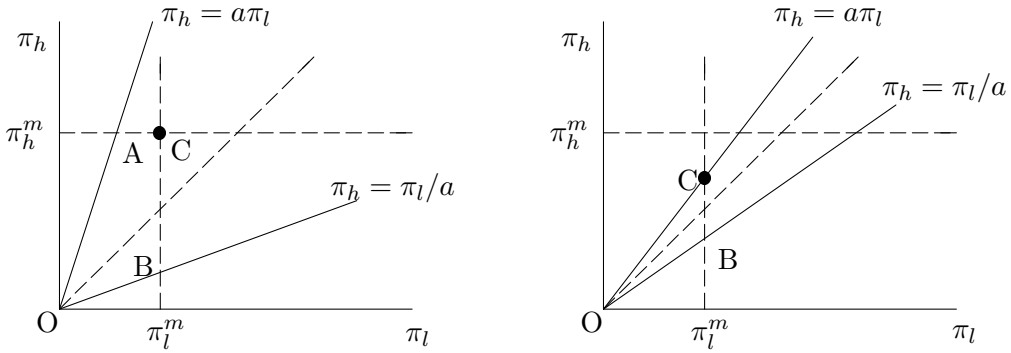

Figure 6.1: Tacit Collusion Under Conditions of Variable Demand. In the left pane the discount rate is sufficiently low that the monopoly price can be sustained in each period (at C), and the right pane is the case where the discount rate is so high that price must be set below the monopoly price in periods of high demand to deter defection.

Notice that in Figure 6.1, the pair of prices (π_l, π_h) must lie to the right of $\pi_l = \pi_l^m$, below $\pi_h = \pi_h^m$, below $\pi_h = a\pi_l$, and above $\pi_h = \pi_l/a$; i.e., in the quadrilateral OACB in the left pane. Clearly total profits are maximized at C, where the monopoly price is charged in both periods. But if a is close to unity (i.e., δ is close to 1/2), the set of feasible prices reduces to the triangle OCB in the right pane, with a maximum again at C. In this case, the firms charge the monopoly price in the low demand state, but something *below* the monopoly price in the high demand state.

6.4 The Folk Theorem: An Embarràs de Richesses

The *Folk Theorem for repeated games* is so called because no one can discover who first thought of it—it is just part of the 'folklore' of game theory. We shall first present a stripped-down analysis of the Folk Theorem with an example, and provide a somewhat more complete discussion in the next section.

Consider the stage game in §6.1. There is of course one subgame perfect Nash equilibrium in which each player gets zero. Moreover neither player can be forced to receive a negative payoff in the repeated game based on this stage game, since zero can be assured simply by playing D. Also, any point in the region OEABCF in Figure 6.2 could be attained in the stage game, assuming the players could agree on a mixed strategy for each. To see this, notice that if Big Fish uses C with probability α and Little Fish uses C with probability β, then the expected payoff to the pair is $(8\beta - 3\alpha, 8\alpha - 3\beta)$, which traces out every point in the quadrilateral OEABCF for $\alpha, \beta \in [0, 1]$ (check it out!). Only the points in OABC are superior to the universal defect equilibrium (0,0), however.

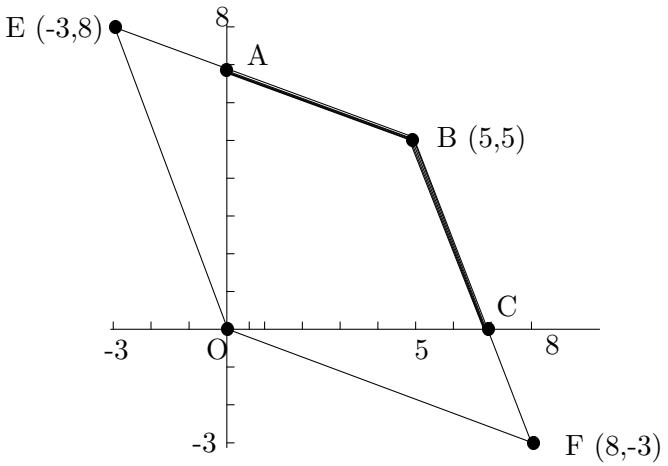

Figure 6.2: The Folk Theorem: any point in the region OABC can be sustained as the subgame perfect Nash equilibrium of the repeated game based on the stage game in §6.1.

The Folk Theorem says that *any point in the region 0ABC can be sustained as a subgame perfect Nash equilibrium of the repeated game based on the stage game in §6.1, provided the discount rates of the players are sufficiently low.*[2] To implement a particular point in OABC, let (α, β)

[2]It is worth repeating that the Folk Theorem is a staple of classical game theory,

be C-probabilities for the two players associated with the point. Allow a random number generator to choose a sequence of choices for Big Fish in which the frequency of C is α, and similarly for Little Fish. Give both players a copy of the list, and instruct both players to play their strategy, and if their partner does not play the appropriate corresponding strategy, to defect (play D) forever. Clearly these strategies are Nash subgame perfect. There is an obvious generalization of this result.

Theorem 6.9. ***The Folk Theorem.*** *Consider any two-player stage game with a Nash equilibrium with payoffs (a, b) to the two players. Suppose there exist a pair strategies for the two players that gives the players (c, d). Then if $c \geq a$ and $d \geq b$, and the discount factors of the players are sufficiently close to unity, there is a subgame perfect Nash equilibrium of the repeated game with expected payoffs (c, d) in each period.*

6.5 Variations on the Folk Theorem

Consider an n-player game with player strategy sets S_i for $i = 1, \ldots, n$, so the set of strategy profiles for the game is $S = \{(s_1, \ldots, s_n) | s_i \in S_i, i = 1, \ldots, n\}$. The payoff for player i is $\pi_i(s)$, where $s \in S$. For any $s \in S$ we write s_{-i} for the vector obtained by dropping the ith component of s, and for any $i = 1, \ldots, n$ we write $(s_i, s_{-i}) = s$. For a given player j, suppose the other players choose strategies m^j_{-j} such that j's best response m^j_j gives j the lowest possible payoff in the game. We call the resulting strategy profile m^j the *maximum punishment payoff* for j. Then $\pi^*_j = \pi_j(m^j)$ is j's payoff when everyone else "gangs up on him." We call $\pi^* = (\pi^*_1, \ldots, \pi^*_n)$ the *minimax point* of the game. Now define

$$\Pi = \{(\pi_1(s), \ldots, \pi_n(s)) | s \in S, \pi_i(s) \geq \pi^*_i, i = 1, \ldots, n\},$$

so Π is the set of strategy profiles in the stage game with payoffs at least as good as the maximum punishment payoff for each player.

We let this be the stage game for a repeated game with discount factor δ, common to all the agents. If the stage game is played in periods $t = 0, 1, 2, \ldots$, and if the sequence of strategy profiles used by the players is $s(1), s(2), \ldots$, then the payoff to player j is

$$\tilde{\pi}_j = \sum_{t=0}^{\infty} \delta^t \pi_j(s(t)),$$

but is widely disrespected in evolutionary game theory, because few of the equilibria it supports are stable with respect to a plausible evolutionary dynamic.

With these definitions behind us, we have

Theorem 6.10. ***The Folk Theorem.*** *For any $\pi = (\pi_1, \ldots, \pi_n) \in \Pi$, if δ is sufficiently close to unity, there is a Nash equilibrium of the repeated game such that π_j is j's payoff for $j = 1, \ldots, n$ in each period.*

The proof of the theorem is straightforward. For any such $\pi \in \Pi$, each player j uses the strategy s_j that gives payoffs π in each period, provided the other players do likewise. If one player deviates, however, all other players play the strategies that impose the maximum punishment payoff on j forever. Since $\pi_j \geq \pi_j^*$, player j cannot gain from deviating from s_j, so the profile of strategies is Nash.

Notice that this Folk Theorem says nothing about subgame perfection. Indeed, the strategies that impose the maximum punishment on player j may be far from a Nash equilibrium, so the punishment threat is not credible. However we have:

Theorem 6.11. ***The Folk Theorem with Subgame Perfection.*** *Suppose $y = (y_1, \ldots, y_n)$ is the vector of payoffs in a Nash equilibrium of the underlying one-shot game, and $\pi \in \Pi$ with $\pi_i \geq y_i$ for $i = 1, \ldots, n$. Then if δ is sufficiently close to unity, there is a subgame perfect Nash equilibrium of the repeated game such that π_j is j's payoff for $j = 1, \ldots, n$ in each period.*

The proof of this theorem is also straightforward. For any such $\pi \in \Pi$, each player j uses the strategy s_j that gives payoffs π in each period, provided the other players do likewise. If one player deviates, however, all players play the strategies that give payoff vector y forever.

This theorem is of course much more plausible than the first Folk Theorem, since it does not involve incredible threats. Fudenberg and Maskin (1986) have shown that subgame perfection is possible even when π does *not* dominate any one-shot Nash equilibrium, but this requires a great deal more coordination among the players, since players must be punished for not punishing defectors, and players must be punished for not punishing players who have not punished when they were supposed to, and so on.

There is no doubt but that the repeated game helps us understand social cooperation, and reveals some of the central factors upon which the capacity of a society to sustain cooperation depends—such factors as the discount rate, the probability that the stage game will be replayed, the expected time between plays, the degree of informational completeness,

and the like. However the existence of vast numbers of subgame perfect Nash equilibria renders the repeated game framework virtually useless. By explaining practically anything, the model in fact explains nothing.

It is tempting simply to reject the Pareto-inefficient equilibria out of hand, after which most of the multiplicity disappears. The problem is that even when there is cooperation in repeated interactions, this cooperation is usually less than perfect and often breaks down. Moreover, while trigger strategies are sometimes plausible (e.g., when a couple breaks up, or a boss fires a worker, or a client switches suppliers), they are often not plausible, because the gains from 'renegotiating' a cooperative solution are often quite large. This leads players to defect, expecting that the breakdown of cooperation can be repaired in the not too distant future—that the 'poor repentant sinner' will be taken back into the fold.

6.6 The One Stage Deviation Principle

Suppose s is a strategy profile for players in a repeated game, so s specifies what move each player makes at each stage of the game, depending on the prior history of moves in the game. We say s satisfies the *One Stage Deviation Principle* if no player can gain by deviating from s, either on or off the equilibrium path of the game tree, in a single stage and conforming to s thereafter. The following theorem is often very useful in analyzing repeated games, since it says that a strategy profile is subgame perfect if it satisfies the one stage deviation principle. Clearly it is a lot simpler to check single stage deviations than arbitrary alternative strategy profiles.

Theorem 6.12. ***The One-stage Deviation Principle.*** *A strategy profile s for a repeated game with positive discount rates, based on a stage game with bounded payoffs, is a subgame perfect Nash equilibrium if and only if it satisfies the one-stage deviation principle.*

For a proof of this theorem (it's not difficult), and some extensions, see Fudenberg and Tirole (1991).

6.7 A Trembling Hand, Repeated Game, Cooperative Equilibrium

Suppose the acts involved in cooperating in §6.1 are in fact complex and demanding, so there is some probability $\epsilon > 0$ that a player will make a mistake, playing D instead of C. Of course trigger strategies

are devastating in such a situation, since with probability one eventually one player will make a mistake, and both will defect for ever after. If it were possible to distinguish mistakes from intentional defection, there would be no difficulty in sustaining cooperation. Suppose, however, that there is not. There may nevertheless be a *trembling hand* cooperative equilibrium of the following form: if either player defects, both defect for k rounds, and then both return to cooperation (no matter what happened in the defection rounds). Given ϵ, of course k should be chosen to be the *smallest* integer such that it pays to cooperate rather than defect.

Does such a k exist? Let v be the value of the game when both players use the following 'trembling hand' strategy. There are two 'phases' to the game. In the 'cooperate' phase, (try to) play C, and in the 'punishment' phase, play D. If the game is in the cooperate phase and either agent plays D (on purpose or by mistake), the game moves to the punishment phase. If the punishment phase has lasted for k rounds, the game moves to the cooperate phase. The game starts in the cooperate phase.

It is clear that there is no gain from playing C in the punishment phase. Can there be a gain from playing D in the cooperate phase?

Here is where the *One Stage Deviation Principle* becomes useful. If the above strategies do not form a subgame perfect Nash equilibrium, then playing D in the cooperate phase and then returning to the trembling hand strategy has a higher payoff than cooperating. The payoff to playing D in the cooperate phase and then returning to the trembling hand strategy is just $8(1-\epsilon) + d^{k+1}v$, since your partner also plays D with probability ϵ, in which case you get nothing but still must wait k periods to resume cooperating. Thus cooperating is Nash when $8(1-\epsilon)+d^{k+1}v$, is less than v, or

$$v(1-\delta^{k+1}) > 8(1-\epsilon). \tag{6.4}$$

But what is v? Well, v must satisfy the equation

$$v = (1-\epsilon)\epsilon(-3+\delta^{k+1}v) + (1-\epsilon)^2(5+\delta v) + \epsilon[(1-\epsilon)8 + \delta^{k+1}v]. \tag{6.5}$$

To see this, note that if you both cooperate at the start of the game, with probability $(1-\epsilon)\epsilon$ you play C and your partner plays D, in which case you get -3 now and after $k+1$ periods, you're back into the cooperate phase, the present value of which is $\delta^{k+1}v$. This is the first term in (6.5). With probability $(1-\epsilon)^2$ you both play C, so you get 5 and v again in the next period. This is the second term. With probability ϵ you play D, in which case you get 8 if your partner plays C, zero if your partner plays D, and in either case, you get v after $k+1$ periods. This is the final term in (6.5).

Solving (6.5) for v, we get

$$v = \frac{5(1-\epsilon)}{1-\delta(1-\epsilon)^2 - \delta^k\epsilon(2-\epsilon)}.$$

Suppose, for instance, $\epsilon = 15\%$ and $\delta = 95\%$. It is easy to check by hand calculation that $v = 85$ if there is no punishment phase ($k = 0$), but the payoff to defecting is $87.55 > v$. For one round of punishment ($k = 1$), $v = 67.27$ but the value of defecting is $8(1-\epsilon)+\delta^2 v = 67.51$, so punishment is still not sufficiently severe. For two rounds of punishment ($k = 2$), $v = 56.14$ and $8(1-\epsilon)+\delta^3 v = 54.93$, so two rounds of punishment is sufficient to sustain cooperation. Notice that if we could assume players were 'honest' (always cooperated), then the value of the game would be 85, but the cost of inducing cooperation when we cannot assume 'honesty' reduces the value to 54.93.

6.8 Death and Discount Rates in Repeated Games

Suppose an agent plays a repeated game in which the payoff in each period is π, the agent's discount rate is ρ, and the probability that the agent dies at the end of each period is $\sigma > 0$. Prove that the present value of the game to the agent is the same as the present value of the game would be if the probability of death were zero and the discount rate were $\rho + \sigma$.

6.9 Manny and Moe

Manny and Moe play the game in the figure an indefinite number of times. The do not discount the future. Show that if the probability p of continuing the game in each period is sufficiently large, then it is Nash for both Manny and Moe to cooperate (play C) in each period. What is the smallest value of p for which this is true?

Moe \ Manny	C	D
C	3,3	0,5
D	5,0	1,1

6.10 The Strategy of an Oil Cartel

Suppose there are two oil producing countries, Iran and Iraq. They both can operate at either of two production levels: 2 or 4 million barrels a day. Depending on their decisions, the total output on the world market will be 4,6, or 8 million barrels a day, and the price per barrel in these

three cases is \$25, \$15, and \$10, respectively. Costs of production are \$2 per barrel for Iran and \$4 per barrel for Iraq.

(a) Draw a normal form matrix of this game and its payoffs;

(b) Show that is a prisoner's dilemma and the equilibrium is for both to produce high output;

(c) Now suppose this game is repeated every day, and both countries agree to cooperate by producing the low output, each one threatening the other with a *trigger strategy*: "if you produce high output, even once, I will produce high output forever." Show cooperation is a Nash equilibrium if the discount rate is sufficiently low. Do you think this threat is *really* credible? Wouldn't one of the countries try to renegotiate the deal after a while? Wouldn't both sides know this, and take it into account? How could they take it into account?

6.11 Tit-for-Tat

Consider a repeated prisoner's dilemma with two players, where the payoff to mutual cooperation is (r, r), the payoff to mutual defection is (p, p), and the payoff to defecting against a cooperator is t for the defector and s for the cooperator, where $t > r > p > s$. A *Tit-for-Tat* player cooperates on the first round of a repeated prisoner's dilemma and thereafter does what the other player did on the previous round.

(a) Show that for a sufficiently low discount rate for each of the players, Tit-for-Tat is a Nash equilibrium.

(b) What is the largest discount rate that can support a Tit-for-Tat Nash equilibrium?

(c) Show that Tit-for-Tat is not subgame perfect. Hint: If a player defects on a particular round, Tit-for-Tat specifies that the two players will exchange cooperate/defect forever, which has a lower payoff than cooperating forever. You might think that it can't possibly matter, since a player using a trigger strategy will *never* be the first to defect. But this is irrelevant—the definition of a strategy in a game requires that a move be chosen for *every node in the game tree where the player moves*, whether or not the node can be reached with positive probability in a Nash equilibrium. In fact, often what happens at unreachable nodes—nodes that are not

hit if players play a particular Nash strategy profile—is critical to specifying the equilibrium. Subgame perfection is a case in point (if you don't understand this, review Chapter 5).

(d) Revise the Tit-for-Tat strategy to make it subgame perfect. Hint: define *contrite Tit-for-Tat* (Boyd 1989) as follows. A player is either in *good standing* or *bad standing*. In period 1 both players are in good standing. A player is in good standing in period $t > 1$ only if in period $t-1$ (a) he cooperated and his partner was in good standing; or (b) he cooperated and he was in bad standing. Otherwise the player is in bad standing in period $t > 1$. Contrite Tit-for-Tat cooperates unless he is in good standing and his partner is in bad standing. Use the one-shot deviation principle to show that contrite Tit-for-Tat is subgame perfect.

6.12 A Public Goods Experiment

A group of ten economics students in a game theory class play the following game. Each student is given \$1, and is handed the following instructions:

> You may anonymously deposit any portion of your \$1 in a 'public account.' Whatever you do not deposit in the public account, you may keep for yourself. The money in the public account will be multiplied by five and shared among the ten of you.

Thus each student gets \$5 if all cooperate (contribute everything to the public good), and each student gets only \$1 if all defect (contribute nothing to the public account).

(a) Show that this is a ten-player prisoner's dilemma;

(b) What are the Nash equilibria to this game if it is repeated a finite number of times?

(c) What are the equilibria if the game is repeated indefinitely with probability $p = 0.99$, assuming the players have discount rate zero?

(d) What do you think really happens when this game is played (see Chapter 11)?

6.13 Reputational Equilibrium

Consider a firm that can produce a good at any quality level $q \in [0, 1]$. If consumers anticipate quality q_a, their demand x is given by

$$x = 4 + 6q_a - p.$$

Suppose the firm knows this demand curve and takes q_a as given, but can set the quality q supplied. The firm has no fixed costs, and the cost of producing one unit of the good of quality q is $2 + 6q^2$.

In each period $t = 1, 2, \ldots$ the firm chooses a quality level q and a price p. Consumers see the price but do not know the quality until they buy the good. Consumers follow a trigger strategy, in which they buy the good in each period in which $q \geq q_a$, but if $q < q_a$ in some period, they never buy from the firm again.

(a) If the firm uses discount factor $\delta = 0.9$, for what levels of q_a is there a viable *reputational equilibrium*; i.e. an equilibrium in which the firm produces quality q_a in each period?

(b) If consumers know the firm's cost function and discount factor, and anticipate quality q_a that maximizes the firm's profits, and if the firm also chooses q to maximize profits, find the equilibrium level q_a and show that this is a reputational equilibrium.

(c) Suppose now the firm chooses q to maximize profits. Show that they would set $q = 0.5$.

6.14 Contingent Renewal Contracts

In §3.17 we criticized the neoclassical model for purporting to explain the economy without strategic interaction, and in §3.18 we sharpened this criticism by questioning the plausibility of complete contracting. In this section we develop some game theoretic mechanisms that substitute for complete contracting.

In many exchanges, including those between (a) employer and employee; (b) lender and borrower; and (c) firm and customer, the agent on one side of the exchange gives money (employer, lender, customer), while the agent on the other side of the exchange gives a promise (employee, borrower, firm). The employee promises to work hard, the borrower promises to repay the loan, and the firm promises to provide high-quality products. Rarely, however, is this promise subject to a contract that can be enforced at reasonably low cost.

Let us call the player who gives money the *principal*, and the player who gives promises the *agent*. In the absence of an enforceable contract, why do agents keep their promises? Of course, some agents are just honest, but there are doubtless enough dishonest people, and enough people who can convince themselves that they are Doing the Right Thing by *not* keeping their promise, that exchange would break down if enforcement were based purely on the integrity of the agents.

Perhaps the threat of suing is sufficient to secure agent compliance. But generally such threats are not credible. The idea of taking an employee to court for not working hard enough is ludicrous. A lender can sue a borrower for non-payment, but if the borrower was imprudent, winning the suit is a Pyrrhic victory—there is not much to collect! A customer can sue a firm for faulty goods, of course, but very few of us have ever done such a thing, and it is not reasonable to suppose that, except in special cases (for instance products that cause personal injury), firms satisfy customers' needs because they are afraid of being taken to court. So why, then, do agents generally keep their promises?

The answer is, of course, that if agents don't keep their promises, principals dump them: employers fire workers who shirk, lenders refuse future loans to borrowers who have defaulted, and customers switch to new suppliers when dissatisfied. All three actions represent *trigger strategies in a repeated game*: the exchange between principal and agent is renewed indefinitely (perhaps with some exogenous probability of dissolution), the principal using the threat of non-renewal to secure compliance. We call these *contingent renewal* exchanges.[3]

Contingent renewal exchanges are among the most prevalent exchanges in market economies. In my humble opinion, the contingent renewal model should be taught to student in Chapter Two of their Introductory Economics text, right after supply and demand. Only the dead weight of tradition prevents it from being so treated. But this, of course, is obvious, so let's move on.

[3]There are also one-shot principal agent relationships, of course. We discuss these extensively in Chapter 14. Such relationships are probably not very important in the case of firm-consumer or employer-employee exchanges, but they are in the case of lender-borrower exchanges, where reputation and equity sharing are often more important than contingent renewal in protecting the interests of lenders.

6.14.1 Costless Contract Enforcement: Achilles' Heel of Neoclassical Economics

At the heart of neoclassical economics is the notion that the price mechanism equates supply and demand. The Fundamental Theorem of Welfare Economics, in particular, says that in a general equilibrium model, under appropriate conditions every competitive equilibrium is Pareto optimal, and every feasible distribution of utility among agents can be attained by a suitable initial distribution of property rights, followed by competitive exchange (§3.17). There have been many critiques of this model, and despite its enormous contribution to economic theory, it is no longer on the cutting edge of research these days.[4] Non-convexity, externalities, multiple equilibria, and the absence of equilibration mechanisms entailing locally stability, as well as other arcane issues, are often mentioned as the problem with the neoclassical model. But the biggest problem, easily trumping the others, is that the model assumes that *agents can write costlessly enforceable contracts for all goods exchanged on markets.* This may be reasonable for standardized goods (for instance, raw materials, basic chemicals, and easily graded agricultural goods) but, as we have seen, is misleading when applied to labor, credit, and consumer goods markets.

6.14.2 Contingent Renewal Markets Do Not Clear in Equilibrium

A *contingent renewal market* is a market in which exchanges between buyers and sellers are regulated by contingent renewal relationships. Since the principal (employer, lender, consumer) in such markets uses a trigger strategy (the threat of non-renewal) to elicit performance from the agent (worker, borrower, firm), the loss of the relationship must be costly to the agent. But if price is set in such markets to equate supply and demand, the cost to an agent of being cut off by the principal is zero, since the agent will secure another position in the next period at the prevailing price. Hence if the principal uses a trigger strategy, there must be a positive probability that there is an excess supply of agents. It follows that *in a Nash equilibrium of a contingent renewal market, there is an excess supply of agents.*

Notice how nicely this conclusion explains some of the most pervasive facts about market economies.

[4] That it is still taught to students as the 'ideal case' from which the real world is an imperfect realization strikes me as scandalous.

(a) **Labor markets**: In the neoclassical model the wage rate adjusts to equate the supply of and the demand for labor. The general condition of labor markets, however, is *excess supply*. Often this takes the form of explicit unemployment, which neoclassical economist develop complex models to explain, using search costs, friction, adaptive expectations, exotic intertemporal elasticities, and the like. Using Occam's razor (always opt for the simplest explanation first), a contingent renewal labor market does the job. There simply cannot be full employment in such models (Gintis 1976, Shapiro and Stiglitz 1984, Bowles 1985, Bowles and Gintis 1993b). Excess supply in labor markets takes the form not only of unemployment, but 'underemployment': workers hold one position but are capable and willing to fill a 'better' position, even at the going wage or a bit below, but cannot secure such a position.

(b) **Credit markets**: In the neoclassical model the interest rate adjusts to equate the supply of and the demand for loans. The general condition of credit markets, however, is excess demand. Why does the interest rate not rise to cut off this excess demand? There are two basic reasons (Stiglitz and Weiss 1981, Stiglitz 1987). First, an increase in the interest rate will drive borrowers who have low-risk, low-expected-return projects out of the market, and increase the expected riskiness of the remaining pool of borrowers. Second, an interest rate increase will induce borrowers to increase the riskiness of their investment projects, thus lowering the lender's expected return.

Since risk-sharing—requiring the borrower to put up a fraction of the equity in a project—is the most widely used and effective means of endogenous contract enforcement in credit markets, it follows that *lending is directed predominantly toward wealthy agents*. This basic fact of life, which seems so perverse from the neoclassical standpoint (loans should be from the wealthy to the non-wealthy), is perfectly comprehensible from the standpoint of models in which contract enforcement is endogenous, even without contingent renewal. Contingent renewal (making available a line of credit, contingent on performance) adds the dimension that a certain subset of non-wealthy borrowers with good projects can get loans, facing the threat of falling into the pool of unemployed 'credit-seekers' should their credit line be terminated.[5]

[5]There is a wonderful joke illustrating this point. Sol and his lawyer Manny are

(c) **Consumer goods markets**: In the neoclassical model, price adjusts until supply and demand are equal. This implies that firms can sell as much as they want, subject to the market price, and choose how much to produce according to cost considerations. Everyday observation tells a different story: firms try to *Sell More Stuff*, and except in their wildest dreams, they can produce with ease however much stuff they manage to sell. Only under the threat of a failing grade can we convince our students that the common-sense view is wrong. Of course there are sophisticated neoclassical models in which firms have 'differentiated products' and 'downward-sloping demand curves,' but this is not really what is happening. What's really happening is that firms do not have their own, private demand curves—they want to Sell More Stuff at the prevailing price, in head-to-head competition with other firms trying to do the same thing. For instance, automobile manufacturers all serve the same array of markets, and buyers frequently shift among firms. They also sell at about the same price, for a given quality product, and can produce as much as they can sell at constant marginal cost. Downward sloping demand curves do not capture this reality.

If the common-sense view is correct, so sellers want to Sell More Stuff in equilibrium, price must exceed marginal cost in equilibrium. How can this be? The simplest explanation is that *where product quality cannot be ensured by explicit contract, goods are in excess supply.* Consumers typically pay a price *in excess* of marginal cost, the implicit threat to switch suppliers if dissatisfied inducing firms to supply high quality products (Klein and Leffler 1981, Gintis 1989).[6]

riding down Avenue of the Americas in New York, and Sol sees a beautiful office building. "Manny," he says, "find out how much they want for that office building." The next day Manny calls Sol and says, "I have good news and bad news. The good news is they only want $27,000,000, and they'll take 25% of this in the form of stock options." "What's the bad news?" asks Sol. "They want a cash down payment of $50,000," Manny replied.

[6]Is this threat credible? We develop a 'rational actor model' below in which the threat is credible because switching suppliers is costless. Even where switching is costly, however, we can plausibly extend our analysis of retaliation in Chapter 11 to cover the case of costly switching.

6.14.3 Money Confers Short-side Power

Power? That's politics, not economics. Or is it? Suppose we say a principal P has power over an agent A if P can impose, or credibly threaten to impose, sanctions on A, but A has no such capacity vis-à-vis P (Bowles and Gintis 1992). This definition is doubtless incomplete and unnuanced, but conforms to standard notions in analytical political theory (Simon 1953, Dahl 1957, Harsanyi 1962). In the neoclassical model there is no power, because all markets clear and contracts are costlessly enforced. In contingent renewal markets, however, principals have power over agents because they can impose costs on agents by terminating them. Since agents are in excess supply, absent collusion, agents can exercise no parallel threat over their principals. It follows that employers have power over employees, lenders have power over borrowers, and consumers have power over the firms from whom they buy. We may call this *short-side power* because it always lies with the transactor on the short side of the market—i.e., the side for which the quantity of desired transactions is the lesser.

So *contingent renewal markets do not clear, and in equilibrium they allocate power to agents located on the short side of the market.*

6.14.4 When Money Talks, People Listen

If we review the cast of characters in our various contingent renewal markets, we find a strong regularity: the principal gives money to the agent, and the principal is on the short side of the market. For instance the employer, the lender, and the consumer hand over money to the worker, the borrower, and the supplying firm and the latter are all short-siders. The reason for this is clear: the money-side of contracts is relatively easy to enforce. As you will see in Chapter 14 (or you can turn to this chapter now), this important regularity is implicit in all principal-agent models, where it is assumed that the principal can make credible promises (to wit, the incentive scheme), while the agent cannot.

The application of the notion that "money talks" is particularly dramatic in the case of consumer goods markets. In neoclassical theory consumer sovereignty means that free markets (under the appropriate conditions) lead to efficient allocations. What the term *really* means in people's lives is that since firms are on the long side of the market (they are quantity constrained) consumers can tell producers how to behave—people are truly sovereign. Probably nowhere in the daily lives of ordinary people do they feel more power, and gain more respect, than

when acting as consumers, constantly pandered to by obsequious suppliers interested in keeping in their good graces—and benefitting from the difference between price and marginal cost.

6.14.5 The Economy is Controlled by the Wealthy

If you ask average college students whether they agree with the statement that "the wealthy control production and investment in the economy," most will probably reply with some quip like, "Does water run downhill?" "Is the Pope Catholic?" or perhaps with something about the excretory behavior of a bear in the woods. Of course the wealthy control the economy. This is not a criticism, and it's not profound. It's just true.

Economists basing their intuitions on neoclassical economics (§3.17), of course, disagree. The wealthy have great *purchasing power*, they opine, but this does not translate into power in any *political* sense. As Paul Samuelson (1957):894 has noted, "in a perfectly competitive market it really doesn't matter who hires whom; so let labor hire capital." The result, expressed long ago by Joseph Schumpeter (1934) is a decentralization of power to consumers: "The people who direct business firms only execute what is prescribed for them by wants." These views taken together imply the touchingly apolitical conception of the competitive economy expressed by Abba Lerner at the head of this chapter. Unfortunately, it is not always plausible, as we illustrate in the next two examples, one from labor markets and the other from consumer goods markets.

6.15 Contingent Renewal Labor Markets

In this section we develop a repeated game between employer and employee, in which the employer pays the employee a wage higher than (the expected value of) his next best alternative, using the threat of termination (a trigger strategy) to induce a high level of effort, in a situation where it is infeasible to write and enforce a contract for labor effort. When all employers behave in this manner, we have a non-clearing market in equilibrium.

Suppose an employer's income per period is $q(e)$, an increasing, concave function of the effort e of an employee. The employee's payoff per period $u = u(w, e)$ is an increasing function of the wage w and a decreasing function of effort e. Effort is known to the employee but only imperfectly observable by the employer. In each period the employer

pays the employee w, the employee chooses effort e, and the employer observes a signal that registers the employee as 'shirking' with probability $f(e)$, where $f'(e) < 0$. If the employee is caught shirking, he is dismissed, and receives a fallback with present value z. Presumably z depends on the value of leisure, the extent of unemployment insurance, the cost of job search, the startup costs in another job, and the present value of the new job. The employer chooses w to maximize profits. The tradeoff the employer faces is that a higher wage costs more, but it increases the cost of dismissal to the employee. The profit-maximizing wage equates the marginal cost to the marginal benefit.

The employee chooses $e = e(w)$ to maximize the discounted present value v of having the job, where the flow of utility per period is $u(w, e)$. Given discount rate ρ and fallback z, the employee's payoff from the repeated game is

$$v = \frac{u(w,e) + [1 - f(e)]v + f(e)z}{1 + \rho}.$$

where the first term in the numerator is the current period utility, assumed for convenience to accrue at the end of the period, and the others measure the expected present value obtainable at the end of the period, the weights being the probability of retaining or losing the position. Simplifying, we get

$$v = \frac{u(w,e) - \rho z}{\rho + f(e)} + z.$$

The term ρz in the numerator is the forgone flow of utility from the fallback, so the numerator is the net flow of utility from the relationship, while $f(e)$ in the denominator is added to the discount rate ρ, reflecting the fact that future returns must be discounted by the probability of their accrual as well as by the rate of time preference.

The employee varies e to maximize v, giving the first order condition

$$\frac{\partial u}{\partial e} - \frac{\partial f}{\partial e}(v - z) = 0, \tag{6.6}$$

which says that the employee increases effort to the point where the marginal disutility of effort is equal to the marginal reduction in the expected loss occasioned by dismissal. Solving (6.6) for e gives us the employee's best response $e(w)$ to the employer's wage offer w.

We assume that the employer can hire any real number n of workers, all of whom have the effort function $e(w)$, so the employer solves

$$\max_{w,n} \ \pi = q(ne(w)) - wn.$$

The first order conditions on n and w give $q'e = w$, and $q'ne' = n$, which together imply

$$\frac{\partial e}{\partial w} = \frac{e}{w}. \tag{6.7}$$

This is the famous *Solow condition* (Solow 1979).

The best response function and (part of) the employer's choice of an optimal enforcement strategy (w^*) are shown in Figure 6.3, which plots effort against the salary. The iso-v function v^* is one of a family of loci of effort levels and salaries that yield identical present values to the employee. Their slope, $-(\partial v/\partial w)/(\partial v/\partial e)$, is the marginal rate of substitution between wage and effort in the employee's objective function. Preferred iso-v loci lie to the right.

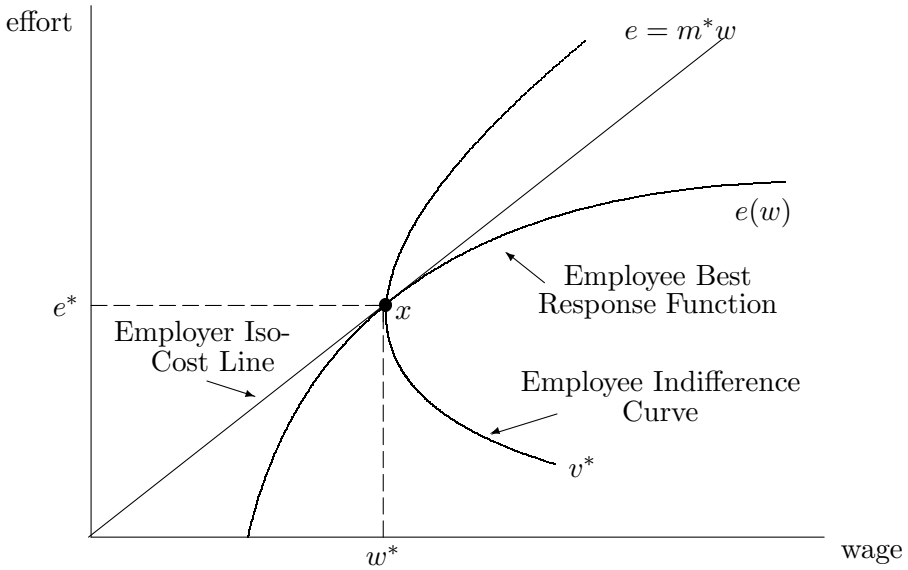

Figure 6.3: The Employee's Best Response Function

By the employee's first order conditions (6.6), the iso-v loci are vertical where they intersect the best response function (because $\partial v/\partial e = 0$). The negative slope of the iso-v functions below $e(w)$ results from the fact that in this region the contribution of an increase in effort (via $(\partial f/\partial e)(v - z)$) to the probability of keeping the job outweigh the effort-disutility effects. Above $e(w)$, the effort-disutility effects predominate. Because v rises along $e(w)$ the employee is unambiguously better off at higher wage. One of the employer's iso-cost loci is labeled $e = m^*w$, where m^* is the profit-maximizing effort per dollar. The employer's

first order condition identifies the equilibrium wage w^* as the tangency between the employer's iso-cost function, $e = m^*w$, and the employee's effort function, with slope e', or point x in the figure.

It should be clear that the contingent renewal equilibrium at x is not first-best, since if the parties could write a contract for effort, any point in the lens-shaped region below the employee's indifference curve v^* and above the employer's iso-cost line $e = m^*w$ makes both parties strictly better off than at x. Notice that if we populated the whole economy with firms like this, we would in general have $v > z$ in market equilibrium, since if $v = z$, (6.6) shows that $\partial u/\partial e = 0$, which is impossible so long as effort is a disutility. This is one instance of the general principle enunciated above, that *contingent renewal markets do not clear in (Nash) equilibrium*, and the agent whose promise is contractible (usually the agent paying money) is on the long side of the market.

Perhaps an example would help visualize this situation. Suppose the utility function is given by

$$u(w, e) = w - \frac{1}{1 - e}$$

and the shirking signal is given by

$$f(e) = 1 - e.$$

You can check that $e(w)$ is then given by

$$e(w) = 1 - a - \sqrt{a^2 + \rho a},$$

where $a = 1/(w - \rho z)$. The reader can check that this function indeed has the proper shape: it is increasing and concave, is zero when $w = 2 + \rho(1 + z)$, and approaches unity with increasing w. The solution for the employer's optimum w, given by the Solow condition (6.7), is very complicated, so I will approximate the solution. Suppose $\rho = 0.05$ and the employment rate is $q \in [0, 1]$. An employee dismissed at the end of the current period therefore has a probability q, of finding a job right away (we assume all firms are alike), and so regains the present value v. With probability $1 - q$, however, the ex-employee remains unemployed for one period, and tries again afterward. Therefore we have

$$z = qv + (1 - q)z/(1 + \rho),$$

assuming the flow of utility from being unemployed (in particular, there is no unemployment insurance) is zero. Solving, we have

$$z = \frac{(1 + \rho)q}{q + \rho} v.$$

For a given unemployment rate q, we can now find the equilibrium values of w, e, v, and z, and hence the employer's unit labor cost e/w. Running this through Mathematica, the equilibrium values of w and e as the employment rate q goes from zero to 0.67 is depicted in Figure 6.4.

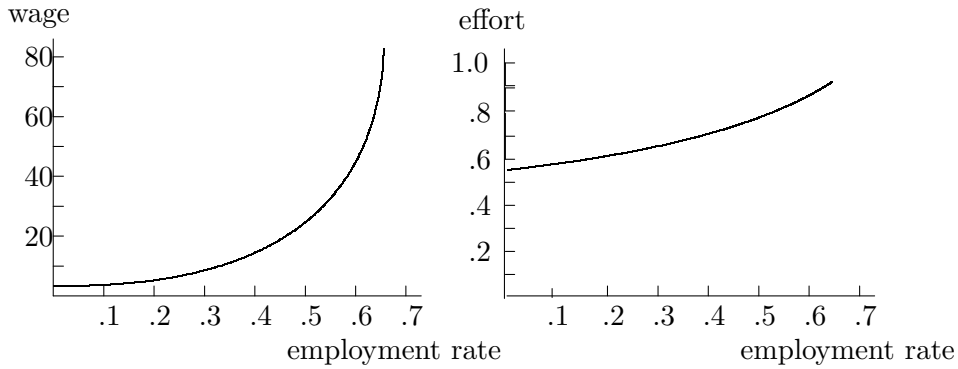

Figure 6.4: Wage and Effort as Functions of the Employment Rate in a Contingent Renewal Labor Market

Notice that while effort increases only moderately as the unemployment rate drops from 100% to 33%, the wage rate increases exponentially as the unemployment rate approaches 33%. I could not find a solution for $q > 0.67$. The actual unemployment rate can be fixed by specifying the firm's production function and imposing a zero profit condition. However this is accomplished, there will be positive unemployment in equilibrium.

6.16 I'd Rather Switch than Fight

In this section we develop a repeated game between firm and consumer, in which product quality cannot be contracted for and can be verified only by consuming the good. The consumer pays a price greater than marginal cost using the threat of brand-switching (a trigger strategy) to induce a high level of quality on the part of the firm. The result is a non-clearing product market, with firms enjoying price greater than marginal cost, and hence are quantity constrained in equilibrium (i.e., they want to Sell More Stuff).

Every Monday families in Pleasant Valley wash clothes. To ensure brightness, they all use bleach. Low quality bleach can, with low but positive probability, ruin clothes, destroy the washing machine's Bleach Delivery Gizmo, and irritate the skin. High quality bleach is therefore

deeply pleasing to Pleasant Valley families. However, high quality bleach is also costly to produce. Why should firms supply high quality?

Since people have different clothes, washing machines, and susceptibility to skin irritation, buyers cannot depend on a supplier's reputation to ascertain quality. Moreover, a firm could fiendishly build up its reputation for delivering high quality bleach, and then when it has a large customer base, supply low quality for one period and then close up shop (this is called "milking your reputation"). Aggrieved families could of course sue the company if they have been hurt by low quality bleach, but such suits are hard to win and very costly to pursue. So no one does this.

If the quality q of bleach supplied by any particular company can only be ascertained after having purchased the product, and if there is no way to be compensated for being harmed by low quality bleach, how can high quality be assured?

Suppose the cost to a firm of producing a gallon of the bleach of quality q is $b(q)$, where $b(0) > 0$ and $b'(q) > 0$ for $q \geq 0$. Each consumer is a customer of a particular supplier, and purchases exactly one gallon of bleach each Friday at price p from this supplier. If dissatisfied, the customer switches to another supplier at zero cost. Suppose the probability of being dissatisfied, and hence of switching, is given by the decreasing function $f(q)$. We assume an infinite time horizon with a fixed discount rate ρ.

(a) Considering both costs $b(q)$ and revenue q as accruing at the end of the period, show that the value $v(q)$ to a firm from having a customer is

$$v(q) = \frac{p - b(q)}{f(q) + \rho}.$$

(b) Suppose the price p is set by market competition, so is exogenous to the firm. Show that the firm chooses quality q so that

$$p = b(q) + b'(q)g(q)$$

where $g(q) = -[f(q) + \rho]/f'(q)$, provided $q > 0$.

(c) Show that quality is an increasing function of price.

Notice that in this case firms are quantity constrained, since price is greater than marginal cost in market (Nash) equilibrium, and that consumers are on the long side of the market. This illustrates the general arguments in §6.14.4.

This model raises an interesting question. What determines firm size? In the standard perfect competition model, firm size is determined by the condition that average costs are at a minimum. This is of course just silly, since a firm can always produce at any multiple of the 'optimal firm size' simply by working the production process, whatever it might be, in parallel.[7] The monopolistic competition model, in which a firm has a downward sloping demand curve, is better, but it doesn't apply to a case like ours, where firms are price takers, as in the perfect competition model, and firm size is determined by the dynamic process of movement of customers among firms. Here is one plausible model of such a process.

Suppose there are n firms in the bleach industry, all selling at the same price p. Suppose firm j has market share m_j^t in period t. Suppose for $j = 1, \ldots, n$, a fraction f_j of firm j's customers leave the firm in each period, and a fraction a_j of customers who have left firms are attracted to firm j. We say the bleach industry is *in equilibrium* if the market share of each firm is constant over time. We have

Theorem 6.13. *There is a unique asymptotically stable equilibrium in the bleach industry.*

Proof: We normalize the number of customers in Pleasant Valley to one. Then the number of customers leaving firm j is $m_j^t f_j$, so the total number of customers looking for new suppliers is $\sum_j m_j^t f_j$. A particular firm j attracts a fraction a_j of these (this assumes a firm can woo back a fraction a_j of its recently-departed customers; the argument is the same if we assume the opposite). Thus the net customer loss of firm j in period t is

$$f_j m_j^t - a_j \sum_{k=1}^{n} f_k m_k^t. \tag{6.8}$$

In equilibrium this quantity must be zero and $m_k^t = m_k$ for all t and for $k = 1, \ldots, n$. This gives the equilibrium condition

$$m_j = \mu_j \sum_{k=1}^{n} f_k m_k, \tag{6.9}$$

[7] Why they teach the standard model to students these days is quite beyond me—it's totally bogus, since its a model of *plant* size, not *firm* size. The important questions of vertical and horizontal integration, the real determinants of firm size, are virtually orthogonal to the question of plant size. Industrial economists have know this for a very long time—for a contemporary review of the literature on the subject, see Sutton (1997).

where we have defined $\mu_k = a_k/f_k$. Notice also that if we add up the n equations in (6.8), we get zero, so $\sum_k m_k^t = 1$ for all t, implying $\sum_k m_k = 1$. Summing (6.9), we arrive at the equilibrium conditions

$$m_j = \frac{\mu_j}{\sum_k \mu_k}.$$

Thus there exists a unique industry equilibrium. To prove asymptotic stability, we define the $n \times n$ matrix $B = (b_{ij})$ where $b_{ij} = a_i f_j$ for $i \neq j$, and $b_{ii} = a_i f_i + (1 - f_i)$, $i, j = 1, \ldots, n$. Then writing the column vector $m^t = (m_1^t, \ldots, m_n^t)$, we have $m^{t+1} = Bm^t$, and hence $m^t = B^t m^0$, where B^t is the t-th power of B. The matrix B is a positive matrix, and it is easy to check that it has eigenvalue 1 with corresponding positive eigenvector $m = (m_1, \ldots, m_n)$. By Perron's Theorem (see for instance Horn and Johnson (1985), Section 8.2), 1 is the unique maximal eigenvalue of B. Also $(1, 1, \ldots, 1)$ is a right eigenvector of B corresponding to the eigenvalue 1. It follows that B^t tends to the matrix whose columns are each m (see Horn and Johnson (1985), Theorem 8.2.8), which proves the theorem. Q.E.D.